

# Are tiled display walls needed for astronomy?

*Bernard F. Meade<sup>A,B,C</sup>, Christopher J. Fluke<sup>A</sup>, Steven Manos<sup>B</sup> and Richard O. Sinnott<sup>B</sup>*

<sup>A</sup> Centre for Astrophysics and Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn, Australia, 3122

<sup>B</sup> The University of Melbourne, Parkville, Victoria, Australia, 3010

<sup>C</sup> Email: bmeade@unimelb.edu.au

**Abstract:** Clustering commodity displays into a Tiled Display Wall (TDW) provides a cost-effective way to create an extremely high resolution display, capable of approaching the image sizes now generated by modern astronomical instruments. Many research institutions have constructed TDWs on the basis that they will improve the scientific outcomes of astronomical imagery. We test this concept by presenting sample images to astronomers and non-astronomers using a standard desktop display (SDD) and a TDW. These samples include standard English words, wide field galaxy surveys and nebulae mosaics from the Hubble telescope. Our experiments show that TDWs provide a better environment than SDDs for searching for small targets in large images. They also show that astronomers tend to be better at searching images for targets than non-astronomers, both groups are generally better when employing physical navigation as opposed to virtual navigation, and that the combination of two non-astronomers using a TDW rivals the experience of a single astronomer. However, there is also a large distribution in aptitude amongst the participants and the nature of the content also plays a significant role in success.

**Keywords:** tiled display wall, optiportal, astronomy, ultra-high resolution images, scalable adaptive graphics environment, sage

## 1 Introduction

Astronomy produces some of the largest volumes of scientific data. Future facilities such as the Large Synoptic Survey Telescope (Tyson 2002; Ivezić et al. 2008) and the Square Kilometer Array (SKA)<sup>1</sup> will produce final datasets heading toward, or even beyond, exabyte sizes.

In this ‘big data’ era of astronomy, existing data analysis tools and methodologies, where the astronomer works directly on the data at the desktop, will be pushed to their limits. There will be an ever-increasing reliance on automated processes to identify objects of interest. This includes the growing variety of generic approaches referred to as data mining (Ball & Brunner 2010; Brescia et al. 2012; Way et al. 2012), and discipline specific solutions such as automated source finders [e.g. see (Koribalski 2012) for a recent review of HI source finding strategies].

As valuable as automatic analyses of these enormous datasets are, astronomy still relies heavily on visual inspection. As the sensitivity of telescopes and detectors is improved, phenomena are increasingly being revealed at the boundary between the signal and the noise. In many cases, these phenomena are not even predicted, making automatic analysis meaningless. It is often a case of not knowing what you are looking for until you see it (Hassan & Fluke 2011).

Not only is the total volume of astronomy data increasing, but the size of individual images (and data

cubes) is growing as well. For example, one of the highest resolution cameras currently available is the Dark Energy Camera (DECam), part of the Dark Energy Survey (DarkEnergySurveyWeb 2012). This camera uses an array of 62 x 2048x4096 CCDs to form a 520 Megapixel image (Mohr et al. 2012). However, as Table 1 demonstrates, there is a growing divide between the resolution of images that can be recorded and the resolution of images that can be displayed on the desktop.

When exploring astronomical imagery, it is desirable to display an image at its native resolution, where there is a one-to-one correspondence between image and display pixels. A very large display with low resolution may reveal less information than a smaller display with a higher resolution. Such high-resolution images reveal more than just the detail of individual celestial objects. In fact, it is the combination of detail and context that make these images valuable: understanding the environment is critical to describing the phenomenon itself. When the image dimensions exceed the capabilities of a standard desktop display, then it is time to look to a non-standard display such as a tiled display wall.

### 1.1 Tiled Display Walls

A tiled display wall is an ultra-high-resolution display comprising a two-dimensional matrix of lower resolution display components, typically standard flat-screen monitors. While there are some slight differences in

<sup>1</sup><http://www.skatelescope.org>

Table 1: Comparison between the typical displays available to an astronomer, and the resolution of some of the current and proposed astronomical cameras. MPs = Megapixels.

Capture Device	Resolution	MPs	Reference
HST Advanced Camera for Surveys	2x2048x4096	16	(ACSweb 2012)
Skymapper	32x2048x4096	268	(Keller et al. 2010)
DECam	62x2048x4096	520	(Mohr et al. 2012)
Subaru Hyper Suprime-Cam	104x2048x4096	870	(HyperSuprimeCamWeb 2012)
Display Device	Resolution	MPs	
Standard desktop display	1680x1050	1.7	
Full high-definition (FHD) desktop display	1920x1200	2.3	
iPad (with retina display)	2048x1536	3.1	
Dell UltraSharp desktop display	2560x1600	4.1	
Laptop (Macbook Pro)	2880x1800	5.2	
4K ultra-high definition (UHD) display	3840x2160	8.3	

the way specific tiled display walls are assembled and configured, e.g. Hiperwalls (HiperwallWeb 2012), Powerwalls (LeedsPowerwallWeb 2009), or OptIPortals (see Appendix A.1), they still operate in a similar manner and hereafter are described simply as TDWs.

A key element of the design principle of TDWs is the use of commodity computers and displays. The computing power available in a standard desktop device with a typical graphics card is capable of driving stunning graphics across multiple displays at very high frame rates. Similarly, the expansion of capabilities of devices such as the emergence of multi-head graphics cards and additional expansion slots on motherboards, means that a single computer can now drive many displays. In fact, a modern computer containing a motherboard with three PCI-Express slots, each hosting a dual-head graphics card, with six Matrox TripleHead2Go<sup>2</sup> devices on each output, can drive 18 full high definition (FHD) displays.

While most TDWs are designed and built as flat screens, either free-standing or mounted on a wall, the use of individual display elements provides a great deal of flexibility in the geometrical configuration. The Mechdyne CAVE2 systems at the Electronics Visualisation Lab (EVL) (University of Chicago) (Febretti et al. 2013) and Monash University, Australia, wrap the TDW around the user, providing an extremely high-resolution immersive stereoscopic environment (74 Megapixels in 2D or 37 Megapixels in 3D). Two key advantages of using monitors over large-screen rear-projection, the usually approach for Cave Automatic Virtual Environments (CAVE; Cruz-Neira et al. 1992), is the increase in both the available pixels and the display brightness. A third advantage is the great reduction in physical footprint of the facility compared to the CAVE, which requires extra space outside of the walls to house the data projectors. The trade-off is a more complex computing and network back-end to drive ~ 80 individual panels, rather than the (maximum) six walls of a cubic CAVE. Additionally, there is the visible presence

of screen bezels - the frame around each of the display elements.

While an ideal TDW would provide a seamless image, in reality the screen bezels introduce a windowing effect. Bezels can be distracting for certain types of content (e.g. office applications), whereas for other tasks they can actually provide a natural coordinate grid to aid in exploration (see Section 3.4). The display panels themselves continue to improve, including the appearance of screens with very thin bezels, such as the Christie Digital FHD551-X<sup>3</sup> with only 5.5mm combined bezel width.

We distinguish between resolution and pixel density when displaying images. For example, the first release of the Retina display for the 15 inch Apple Macbook Pro<sup>4</sup> had a resolution of 2880x1800, which greatly exceeds the typical resolution of a FHD home theatre display at 1920x1080. However, the home theatre's 2 megapixel display can extend over 150 inches (measured diagonally), while the Macbook Pro display crams its 5.1 megapixels into a 15 inch screen. The pixel densities of each configuration are at the extremes, with the Macbook Pro Retina providing a practically seamless image, while the FHD image projected to 200 inches would reveal the individual pixels quite clearly. At this time, there remains a significant price jump to move from FHD to the next off-the-shelf resolution of 2560x1600 pixels (e.g. Dell Ultrasharp). However, the recent emergence of commercially available 4K systems and the Retina displays from companies like Panasonic and Apple, will likely drive down the price of the 2560x1600 displays.

A number of applications exist to simplify the management of a TDW. The two main contenders are Scalable Adaptive Graphics Environment (SAGE; SAGEWeb 2013) from the University of Chicago's EVL and the Cross-platform Cluster Graphics Library (CGLX; CGLXWeb

<sup>2</sup><http://www.matrox.com/graphics/en/products/gxm/>

<sup>3</sup><http://www.christiedigital.com/en-us/digital-signage/products/lcd-flat-panels/pages/55-hd-lcd-flat-panel.aspx>

<sup>4</sup><http://www.apple.com/au/macbook-pro/features-retina/>, 2013



Figure 1: The 29566 x 14321 pixel Carina Nebula mosaic from Hubblesite.org, with OziPortal (15360x6400), Dell Ultrasharp (2560x1600) and Standard Desktop Display (1680x1050) sizes overlaid.

2012; Ponto et al. 2011) from UC San Diego’s CALIT2, though several other solutions also exist. One of the principle benefits of SAGE is that it makes sharing content between TDWs easy (Fujiwara et al. 2011). More recently, Tada et al. (2011) have developed a visualization adaptor that extends the capability of SAGE to allow the display of any X-Window, which opens up the possibility of using almost any application on the TDW.

## 1.2 Background

Astronomical imagery is often seen on promotional material for TDWs, such as the SAGE gallery images<sup>5</sup>. Indeed, the very high resolution images captured by modern detectors (Table 1) do seem very well suited to the environment. TDWs have been used successfully as public outreach devices, such as the displays at the National Institute of Information and Communications Technology, Japan (Morikawa et al. 2010; NICTWeb 2012), and the Adler Planetarium, USA (Adlerweb 2012). However, there is a paucity of literature examining whether TDWs actually do improve understanding of any ultra-high resolution image in any scientific discipline.

Ball & North (2005a) conducted some of the first experiments to compare individual computer dis-

plays with TDWs. They tested subjects on a single desktop display (17 inch, 1280x1024 pixels), a 2x2-display TDW (2560x2048 pixels) and a 3x3-display TDW (3840x3072 pixels). For a target search (red shapes on a black background with random grey dots), they found that participants performed far better when searching for small targets when they could see all the targets at once. The first part of the experiment involved the participants searching for a specific configuration of red dots, while the second part required the participants to match pairs of configurations. No statistically significant difference was observed for targets that could be easily seen on all display environments without needing to pan or zoom.

Ball & North (2005a) also found that the experience of virtual navigation, that is, using a mouse to zoom and pan on a single display, caused more frustration for the participants than the physical navigation required for the TDW. Here, physical navigation means observing parts of an image by physically moving the eyes, head or whole body to an optimum position. The more virtual navigation that was required, the greater the disorientation and agitation experienced by the participants. The authors suggest that being able to easily maintain context while searching for detail made the TDW a more acceptable experience. They also found that the physical construction of the TDW with the screen bezels dividing the image into segments aided the search process. We comment on this issue in Section 3.4.

<sup>5</sup><http://http://www.sagecommons.org/community/sage-walls/>, 2013

Yost et al. (2007) studied the visual acuity of human perception with regards to high-resolution displays. In this context, visual acuity refers to the ability to perceive all displayed information: when an entire display surface is within the user's field of view, and the individual pixels remain discernible, the display is said to be within visual acuity. While increasing the pixel density is one way to exceed visual acuity, increasing both the pixel count and the display size is another. When a display exceeds visual acuity, there are always pixels that cannot be accurately perceived without physical navigation. However, increasing physical navigation does not negatively impact on the performance on most tasks, whereas virtual navigation of the same data or image does have a significant negative impact. This is contrary to the notion that there is no value in using a display that exceeds visual acuity. For practical reasons, it is harder to establish at what point this advantage disappears.

Andrews et al. (2011) focused their investigation on the human experience of using high-resolution displays. In this study, they defined large displays as being human-scale, that is, where the physical size of the display was of similar height and width to naturally occupy the natural field of view of an adult human. In this definition of a large, high-resolution display, this can be achieved through tiling of displays or using individual displays with greater pixel count and large physical size. The authors argue that the design of such displays as TDWs would greatly benefit from considering the physical nature of human-centric search techniques and creating displays that meet these needs.

Andrews et al. (2011) also considered the natural perceptions of users and how a display can affect these perceptions. There is a potential for TDWs to overwhelm the user with information, or induce physical fatigue due to the increased requirements of physical navigation of the environment. However other studies have shown that physical navigation can outperform virtual navigation for many tasks (Ball et al. 2007), and that users quickly adjust to the information density shown on a large display (Andrews et al. 2010). The increased physical activity required for a large display has not shown significant increase in fatigue of subjects, though there is a possibility of some discomfort in the neck due to the increased turning of the head (Ball & North 2005b; Bi & Balakrishnan 2009).

Following on from these design questions, Bezerianos & Isenberg (2012) conducted experiments to determine how proximity to a TDW could affect the perceptions of the user, particularly focusing on angular distortion effects. The ability of participants to effectively estimate quantities such as angle, area and length of objects within an image were significantly affected when the angle of presentation was increased. Thus when a participant was very close to the TDW, their ability to estimate these quantities diminished as the distance of the object to the subject increased. Of particular interest are the results from the second experiment where some of the participants were required to remain in a fixed location, while others were allowed to move freely. The study found that the static position yielded just as accurate results of the mobile

position, but was less time consuming. Therefore, the authors' recommendation that users be encouraged to remain at a distance from the TDW where possible, or physically inspect objects positioned close to their position, ties in closely with the ball2005a observation that users naturally avoid using virtual navigation unless they absolutely have to (Ball et al. 2007).

Most of these studies focus on generalized examples of use of TDW, but the nature of the research disciplines also needs to be considered when investigating these displays. For example, the way an astronomer would use a TDW could have significant differences to the way an economist would use it. As Moreland (2012) argues, we already know how to build the displays, but we have little experience in considering domain-specific applications. The desirability of achieving a one-to-one correspondence between image and display pixels aside, it is far too simplistic to suggest that images of  $A \times B$  pixels require displays of equivalent resolution. Instead, the need must be borne out of the research and the data, where the impact of virtual navigation impedes comprehension.

### 1.3 Overview

In this paper, we describe a series of experiments designed to investigate the assumption that TDWs are intrinsically beneficial in astronomical research. We focus our attention on targeted searches within high resolution images that exceed the available resolution of a standard desktop display. We consider the performance of both individuals and pairs of users at finding targets of decreasing size on either a standard desktop display or a TDW. The participants in the experiments included professional astronomers, experienced amateur astronomers and non-astronomers.

The TDW used in these experiments, the OzIPortal, was built by the School of Engineering at the University of Melbourne in 2008 and is now operated by the University's central IT department. The TDW comprises a 6x4 matrix of Dell Ultrasharp monitors (2560x1600). With a total resolution of 15360x6400 pixels, it is capable of displaying 98.3 Megapixels. However, as Figure 1 shows, this is less than a third of the pixels available in images such as Hubble's Carina Nebula mosaic<sup>6</sup>.

The OzIPortal initially used CGLX for the interface, but this was replaced with the somewhat more versatile SAGE software. We describe the history of the OzIPortal in Appendix A.

The remainder of this paper is set out as follows. In Section 2, we describe the OzIPortal experiments, including the image selections, participants and procedure. In Section 3, we show the experimental results. We look at the comparative performance of targeted searches using both standard desktop display and TDW environments. We consider the performance of the non-astronomer, astronomer and collaborative pair groups. We comment on key findings from the post experiment survey and video observations. In

<sup>6</sup>[http://imgsrc.hubblesite.org/hu/db/images/hs-2007-16-a-full\\\_jpg.jpg](http://imgsrc.hubblesite.org/hu/db/images/hs-2007-16-a-full\_jpg.jpg)



Figure 2: Survey results for self-rated level of expertise with astronomical imagery, for the astronomer and non-astronomer groups. The non-astronomer cohort (green) self-identified strongly with the low end of the experience scale, while the astronomer group (red) is towards the high end.

Section 4 we discuss the implications of these results in the context of the potential use of TDWs in astronomy. We consider further experiments that are either extensions of the current work, or alternative aspects of using a TDW that might be beneficial to astronomers. Concluding remarks are made in Section 5.

## 2 The OzIPortal experiments

In this section, we describe our experimental procedure to investigate the role TDWs might play in aiding knowledge discovery and comprehension of ultra-high resolution images (i.e.  $\sim 100$  megapixels). These images provide researchers with an opportunity to seamlessly explore both context and detail at will. Yet on a standard desktop display (SDD), defined for our purposes as a 24 inch LCD with a resolution of 1680 x 1050 pixels<sup>7</sup>, a researcher must choose dynamically between context or detail, as both cannot be seen at once. In particular, we wanted to determine if there was indeed a definable performance improvement when using a TDW compared to a SDD, which corresponds to the popular expectation that big images need a big display to be seen “properly”. The high-resolution images and target objects were chosen from three different categories: English words, galaxies and nebulae (see section 2.2).

### 2.1 Participant selection

Participants were recruited from two different demographic categories: astronomers and non-astronomers.

<sup>7</sup>The recommended size for centrally deployed computers at the University of Melbourne at the time of the experimental work

For the astronomers, participants included academics, postdoctoral researchers, research students and advanced amateurs. Within this group there was a mix of radio, optical and theoretical astronomers. Non-astronomer participants had a wide range of experience with astronomical imagery, ranging from none to a high level of familiarity. As such a secondary category of expert and non-expert was introduced, based on the participant’s self-rated level of experience with astronomical imagery. Figure 2 shows that the non-astronomer cohort self-identified strongly with the low end of the experience scale, while the astronomer group is towards the high end. This self-rating reflects that, for example, a theoretical astronomer may not feel they have the same expertise as an optical astronomer who works constantly with images.

A total of 45 non-astronomers and 12 astronomers participated in a range of experiments. All participants had a reasonably high-level of familiarity with graphical user interfaces and the use of a mouse for panning and zooming, but few had any prior exposure to a TDW. We report here on the performance results of a subset of 30 participants, noting that:

- The first five non-astronomer subjects participated in an experiment refinement phase and thus their performance results have been excluded from the results described below.
- 14 non-astronomers were presented with a slightly revised set of tasks to those described here. These additional tasks focused on a small target search and multiple image inspection. The small target search proved too difficult to complete in the SDD environment due to a “too-restrictive” time limit of two minutes, and too few participants were available to complete the multiple image inspection.

As all of these participants did complete the post-experiment survey, providing relevant comments on issues such as the suitability of the TDW for the target search task, we retained their survey responses for subsequent qualitative interpretation.

16 of the remaining non-astronomers completed a target search in pairs in order to investigate the process of collaborative inspection on SDDs and TDWs (see Section 2.4).

### 2.2 Image and target selection

In order to establish a common ground between the astronomer and non-astronomer groups, the first image was made up of black words on a white background at a resolution that precisely matched the TDW (15360 x 6400), as can be seen in Figure 3. At this resolution, all words were readable on the TDW without the need to zoom the image.

The words were taken from a list of the top 250 English words (AnglikWeb 2003) to ensure all participants were familiar with the targets. The words were rendered in Arial font and were sized in points of 1000, 300, 100, 30, and 10. Five targets were created at each size except the smallest where an extra five were added. For each size, an equivalent number





Figure 3: The OzIPortal TDW with English word targets displayed at their native resolution. Arial font sizes used were 1000, 300, 100, 30, and 10 points. All words were visible on the TDW using physical navigation and no zooming.

Table 2: Images used for the galaxy search and nebula search. Note that the Carina Nebula image was displayed at 50% of the native resolution for performance reasons

Image ID	Field description	Resolution	Targets
Galaxy Set A	The Coma Cluster	10816 × 7679	Figure 4
Galaxy Set B	CANDELS Ultra Deep Survey	15516 × 8255	Figure 5
Nebula Set A	The Carina Nebula (NGC3772)	29566 × 14321 @ 50%	Figure 6
Nebula Set B	HST-Spitzer Composite of Galactic Center	12203 × 4731	Figure 7

of non-target words were added from the same list, to reduce the possibility of participants guessing based purely on size.

When viewed on the SDD, scaling the image to full screen reduced readability to words in 100pt font or greater. For words 30pt or 10pt in size, zooming the image was necessary.

The astronomy targets were chosen to present a range of sizes similar to the word sizes described above, chosen from amongst the largest available on the HubbleSite gallery<sup>8</sup> - see Table 2 for details. For performance reasons, the Carina Nebula mosaic was shown at 50% of the native resolution.

From these images, the search targets were selected to roughly correspond to the physical sizes of the words, without following a strict sizing scale. The largest astronomy target was 2100 × 1730 pixels while the smallest target was 185 × 145 pixels. Images were not ro-

tated, but were scaled to appear the same size on the search target presentation screen. Astronomical targets were chosen to reflect increasing difficulty. Due to the increased difficulty of the astronomical search compared to the word search, and the limited amount of time available for each participant to complete each task, the number of targets was restricted to 10 per image.

Targets selected from the galaxy images included structures around the galaxy. However, these targets exist on a black background and have no visible connectivity to the other objects in the image. Nebulae provide a fully connected structure with details visibly connected to the context. Figures 4 to 7 show each of the astronomical images and the targets. Additional galaxy and nebula images were used to introduce the environments but were not used during the experiment.

In order to eliminate any potential presentation bias, the image sets were shown alternating for the

<sup>8</sup><http://hubblesite.org/gallery/album/entire/hires>true/>




Figure 4: Galaxy Set A targets in the Coma Cluster ( <http://hubblesite.org/newscenter/archive/releases/2008/24/image/a/>)

environments, so that half the participants saw the set A images on the SDD and set B images on the TDW, and vice-versa for the rest.

## 2.3 Procedure

Figure 8 shows the experimental set-up. The individual target objects were presented on a 40 inch LCD TV immediately adjacent to the TDW, as well as on a laptop sitting adjacent to the SDD. A standard Microsoft PowerPoint presentation was used to display the targets to the participant.

As the participant identified the target (or elected to pass on finding a given target), the presentation was advanced to the next target. Participants were given a total of 2 minutes to find as many of the targets as they could. Once the experiment had been completed on the SDD, the participant was then shown a new set of images and targets on the TDW, again given 2 minutes for each set. Several of the experiments were also filmed for later investigation as to how the displays were used.

Participants were introduced to the experiments, with a brief explanation of their purpose and a demonstration of how to use the two types of displays. The SDD was a familiar environment for all participants

and very little introduction to the environment was required. In the case of the SDD, participants were advised that they would be able to find most of the large targets without using the mouse to pan and zoom, but would need to use virtual navigation for the very small targets. The mouse operation was already second nature, though most participants attempted to minimize the mouse use, preferring to lean closer to the screen.

Very few of the participants had ever seen a TDW before and so the experience was entirely new to them. Those that had encountered such a display before showed little if any advantage when engaged in the structured search experiment. The only significant advantage pre-exposure was that ability to “zoom” by physically approaching the TDW was already known.

In the initial experimental refinement phase, a test group of five non-astronomers was given the same introduction to the TDW as they were to the SDD. The result was that these participants all felt obliged to use the TDW in exactly the same way they had used the SDD, i.e. they sat well back from the screen to obtain the same field of view and used a mouse to zoom and pan rather than walk closer to the screen.

Due to the nature of the TDW software, zooming and panning resulted in some slight image tearing as the screen refresh was not always perfectly synchro-




Figure 5: Galaxy Set B targets in the CANDELS Ultra Deep Survey (<http://hubblesite.org/gallery/album/entire/pr2013011b/hires/true/>)






Figure 6: Nebula Set A targets in the Carina Nebula (<http://hubblesite.org/gallery/album/nebula/pr2007016a/hires/true/>)




Figure 7: Nebula Set B targets in the HST-Spitzer Galactic Center composite (<http://hubblesite.org/newscenter/archive/releases/2009/02/image/d/>)




Figure 8: Experiment layout as described in Section 2.3. The individual target objects were presented on a 40 inch LCD TV (Target Display) immediately adjacent to the TDW, as well as on a laptop sitting adjacent to the SDD (Laptop). A standard Microsoft PowerPoint presentation was used to display the targets to the participant. The six columns of the OzIPortal are driven by six column display nodes, with master control under the SAGE environment from the Head node [Image: Carina nebula mosaic from <http://www.hubblesite.org>].

nized. Moreover, the zoom was not visually active with the image jumping between zoom levels rather than scaling dynamically, as participants were used to on their SDD.

While these issues can be mitigated with higher networking speeds, a simple alternative was found: the participants were told that standing and approaching the screen would more effectively function as zoom (i.e. physical navigation). This very simple training was included in the familiarization stage for the later participants. The use of physical navigation greatly improved the user satisfaction and performance with the TDW, and presented a more realistic assessment as to how the displays should be used in practice.

For the non-astronomer singles group, all participants began the target search using the SDD. For the astronomer group and the collaborative groups, we tested several participants with the TDW display first to see if there was any advantage to the order of exposure.

As Figure 10 shows, there is no significant difference in performance on the target searches regardless of the order of the environments used. The galaxy and nebula targets were alternated for the two environments in order to ensure that no advantage could be ascribed to a particular image/environment combination. After completing the experiment in both environments, the participants were then asked to complete a survey about their experience.

## 2.4 Collaborative pairs

16 non-astronomers were asked to complete the experiment in pairs. They received the same introduction as

all the other participants, but no specific instruction was given to guide how they should share the task. They were required to determine the best way to operate between themselves as part of the task, and in all cases settled the matter of who would operate the interface (in the case of the SDD) or how they would split the search area (in the case of the TDW), with a very brief discussion. This process was occasionally completed during the introduction and so took no time during the task, however, in all cases it did not delay the image inspection process as the pairs began searching while discussing. No disadvantage was observed and therefore no adjustment has been made to the performance results. The only practical difference to the conditions of the experiment was that agreement was required from both participants for any ambiguous situation, for example, some targets could be mistaken for a similar looking non-target. This situation included when no target could be found, at which time both participants could agree to “pass”.

## 3 Results

The results of the experiments are presented in the next two sections: section 3.1 shows the empirical results for the different image identification tasks and 3.2 contains our analysis of the post-experiment survey. We report also on the outcomes of the video observations (section 3.3), which provide some valuable clues as to how to make better use of TDWs. In addition, we look at the specific feedback made with regards to the TDW bezels (section 3.4).



Table 3: Search success rates for non-astronomers (10 participants), astronomers (12 participants), non-astronomer collaborations (8 x 2 participants) and a combination of all three (10 + 12 + 8 = 30 sets of results). These results are based on the median values for the word, galaxy and nebula feature searches, with a quoted range of one standard deviation.

	Group	N <sub>P</sub>	SDD	TDW
<b>Word search</b>	Non-Astronomer	10	66% ± 2.0	66% ± 4.4
	Astronomer	12	66% ± 1.8	80% ± 2.6
	Non-Expert (self-rated)	15	68% ± 1.7	72% ± 4.3
	Expert (self-rated)	7	64% ± 2.3	76% ± 3.0
	Collaboration	8 groups	66% ± 3.0	76% ± 5.4
<b>Galaxy search</b>	Non-Astronomer	10	60% ± 1.7	75% ± 2.5
	Astronomer	12	90% ± 1.9	90% ± 1.3
	Non-Expert (self-rated)	15	70% ± 1.8	80% ± 2.4
	Expert (self-rated)	7	90% ± 2.1	80% ± 1.3
	Collaboration	8 groups	80% ± 0.9	100% ± 0.0
<b>Nebula search</b>	Non-Astronomer	10	50% ± 2.0	60% ± 1.3
	Astronomer	12	65% ± 1.2	80% ± 1.6
	Non-Expert (self-rated)	15	60% ± 1.9	60% ± 1.9
	Expert (self-rated)	7	70% ± 1.3	80% ± 1.3
	Collaboration	8 groups	70% ± 1.2	80% ± 1.6
<b>Combined</b>	Non-Astronomer	10	61% ± 3.3	67% ± 5.2
	Astronomer	12	71% ± 2.8	82% ± 3.3
	Non-Expert (self-rated)	15	67% ± 3.1	71% ± 5.3
	Expert (self-rated)	7	71% ± 3.4	78% ± 3.5
	Collaboration	8 groups	70% ± 3.4	82% ± 5.6

### 3.1 Success rates

Figures 11 and 12 show the individual targets successfully identified. Figure 11 shows the first 15 target words were found by most participants. This corresponds to point sizes of 1000, 300 and 100, which were easily readable in both environments. At point size 30, the words were no longer readable on the SDD and therefore virtual navigation was necessary, causing a performance decline. This can be seen by the rapid drop in the SDD success. Very few of the 10 point words were found in the SDD environment. However, the TDW success rate shows only a slight decline for the 30 point words and a steady decline for the 10 point words.

Figure 12 shows that the galaxy images in set A were fairly well matched for success in both environments, while set B showed a higher success rate for the TDW targets. Similarly for the nebula sets, with the exception of target 9 in the Nebula Set A, which was not found in either environment. As can be seen in Figure 6, target 9 was a subset of target 2, but not particularly more difficult than target 9 of Figure 7, also within target 2 for that set.

Table 3 shows the combined success rate for each group in each environment, where success refers to the number of targets identified during the test. These results indicate that generally performance on the TDW is slightly better than for the SDD for the same set of tasks, with the notable exception that self-rated experts actually performed slightly worse on the TDW

for the Galaxy search. However there are other factors to consider. For example, the sample size is fairly small and the task is not necessarily indicative of typical astronomy work.

Table 3 also shows that the attempt to establish a consistent baseline between the cohorts was effective. In the word search on a SDD, where little experiential value could be ascribed, all groups achieved very similar results. Here the targets and the navigation method were familiar to all subjects. However, the non-astronomers did not experience an improvement in performance when searching for words on the TDW. This reason for this is uncertain, but could be because the astronomers' familiarity with exploring large images translated more easily to the TDW environment. Video observations reveal that astronomers tended to adopt methodical search strategies and were quicker to adapt their strategies to the TDW environment than non-astronomers. See Section 3.3 for more information on video observations.

A useful way to view these results is to consider the comparison of the results for the specific environments. Subtracting the results for the SDD from the TDW produces a simple comparison of the two environments, as seen in Figure 9 for the 10 non-astronomers (top panel), 12 astronomers (middle panel) and eight collaborative pairs of non-astronomers (bottom panel). The astronomer group test results are comparable to the non-astronomer cohort. The astronomers did demonstrate slightly better performance overall, particularly with the word and nebula search. However, the galaxy




Figure 9: Comparison of the success rates of the TDW and SDD for 10 non-astronomers (top panel), 12 astronomers (middle panel), and eight collaborative pairs of non-astronomers (bottom panel). Positive numbers show success favoured the TDW, negative results favoured the SDD. The three types of experiments are the word search (green), galaxy search (red) and nebula search (purple).




Figure 10: Results based on the presentation order of the display environments. Each pair of labelled bars indicates the image type and the first environment participants were exposed to (SDD or TDW in brackets). The green and red bars indicate the number of targets found using the SDD and TDW respectively. There is no strong dependence on which display technology that participants used first.

search results show that the astronomers tended to perform equally well on both the SDD and the TDW (c.f. Table 3). Observations supported by the video recordings show that the astronomers tended to have a more systematic approach to searching, and were less confused by targets split by the screen bezels (edges).

The TDW has often been cited as an ideal environment for research collaboration. The bottom panel of Figure 9 shows the results obtained by pairing two non-astronomers for the same task. Table 3 shows that non-astronomer collaborators match or exceed the performance of a single astronomer and show marked improvement of TDW over SDD.

Figure 9 shows an interesting anomaly with one participant finding the SDD to be far better for the word search than the TDW. In this case, the participant overlooked a 1000pt word and began to search among the smaller words. This highlights a potential trap with a TDW in that a large object presented on such a display may be too large to see, with participant approaching the TDW and effectively eliminating their chance of recognizing the word. This may be in part due to the way the brain recognizes words as a whole and therefore may not apply to astronomical targets.

### 3.2 Survey responses

The participants were asked to complete a short survey after the experiment, designed to gauge their experience using the two display environments. Participants were asked to rate their own experience with astronomical imagery, as can be seen in Figure 2.






Figure 12: Success rates based on individual search targets with galaxy and nebula images. The green columns indicate the targets actually found using the SDD and the red columns are for the TDW. Only the results from the astronomer group were used as the image context was not recorded for the non-astronomer group.




Figure 11: Success rate based on individual words. The green columns indicate the target words actually found using the SDD and the red columns are for the TDW. Only the results from the astronomer group were used, in order to align with results shown in Figure 12 as the image context was not recorded for the non-astronomer group.

Figure 13 show the survey results for the Ease of Use of the SDD and the TDW respectively. Results for both astronomers and non-astronomers indicate that the SDD is generally perceived as difficult for this kind of search while the TDW is generally perceived as easy to use for the same.

Participants were also asked to rate the suitability of the two environments to the tasks presented. Figure 14 show that the results for the SDD are skewed toward Unsuitable for the search tasks while the participants found the TDW was generally well suited.

### 3.3 Video observations

Several of the participants were also filmed to record the manner in which they used the display environments.

When using the TDW, several participants found themselves overlooking extremely simply targets, particularly in the word search, by assuming that the target they were seeking must be smaller than it actually was. However, generally the approach to searching was fairly uniform, with the subjects standing back to get an overview of the image, and then approaching promising regions of the image. In the case of words, the advantage to the TDW over the SDD was that even the very smallest of words could be clearly read when close, while the same target on the SDD was not even visible when zoomed to full extents.

However, the experiment was designed to make the first targets easy to find and subsequent targets were

made progressively harder. This gave a distinct advantage to the SDD for the early targets, and therefore considerably more time was available for finding smaller targets. On the TDW, however, the larger targets were sometimes overlooked, occasionally due to the splitting of the target by the screen bezels, or due to the participants' assumption that the target must be smaller than it actually was.

In general however, the video review shows participants were more methodical in searching for small targets on the TDW than on the SDD. It appears that the participants were more easily able to identify individual screens that they had already searched; compared to trying to remember which region they had searched of the image on the SDD. The screens of the TDW made for a simple segmentation that was easy to remember.

Furthermore, participants working with a partner found the TDW to be naturally separated into halves with each partner being responsible for their own half. Working together on the SDD, these participants found ceding control to someone else to be frustrating. Some chose to share the task of controlling the mouse, alternating between tasks, while others simply directed their partner by pointing in the direction they wished to explore. This resulted in some confusion, though some collaborations quickly settled into very effective teamwork.

On the TDW, splitting the screens into left and right did not always produce harmony. When one participant became convinced that the target was not on their side, they began to encroach on their partner's domain. For some, this resulted in an unspoken agreement to swap sides, while not so for others. However, the success of the collaborations between non-astronomers produced results that were generally better than a single astronomer (c.f. Table 3). Unfortunately there were not enough astronomers to test collaborative behaviours.

### 3.4 Feedback on bezels

Participants were also asked to comment on the structure of the TDW and how the "screen elements" impacted on the subjects' ability to complete the task. We defined screen elements as anything that interrupted the subjects' view of the search image. This included the SAGE icons and toolbars and the screen bezels. The results can be seen in Figure 15.

For the non-astronomer group, 19% found the screen elements were a continuous distraction while 25% found them initially distracting but quickly learned to ignore them. 27% found the screen elements did not distract them at all during the tasks. In fact, 19% of the non-astronomer respondents felt the bezels actually aided their search, compared to 10% who felt they made the task more difficult.

The astronomer group found the screen elements to negatively impact on their experience more so than the non-astronomer group. 33% found the screen elements distracting throughout the experiment, while 6% found them initially distracting, but not so later. However, much like the non-astronomer group, 22%

found the screen elements did not distract them from searching for their targets and 17% found the bezels aided their search strategy. However, 22% found the screen elements hindered their efforts.

The general comments relating to the display environments and the screen bezels generally support the results described above. For the SDD, the comments suggest that the environment was well suited to examining very large images in a broad context, when the entire image could be seen. However, small details within these large images were much harder to find and the context was often lost, making strategic searches harder.

Contrasting this with the TDW, the expanse of the display itself sometimes made observing the whole image harder as the participant needed to be much further back to achieve the same field of view, or turning the head considerably more. However, the dynamic nature of physically approaching the screen and being able to see extremely fine detail within the large image made the search for small targets easier on the TDW. This combined with the physical break in the image due to the screen bezels provided subjects with an easier search methodology.

## 4 Discussion

Previous experiments have shown that for searches of small targets within much larger images, a distinct advantage exists for a TDW (Ball & North 2005b; Ball et al. 2007; Yost et al. 2007), when the case for physical navigation being preferable to virtual navigation is clear. However, when the target size varies, the advantage is less apparent, though the case for physical versus virtual navigation remains. This is because the larger targets can be found with relative ease in a large image, even when it is presented subsampled on a SDD.

Our study showed that participants typically attempted to gain an overview of the whole image to identify regions of interest. In the case of a large target, this was often more readily found on the SDD as it was quicker to obtain this overview and therefore ascertain the target. As we chose targets that would not easily be confused with other objects, this continued to be the case as the targets got smaller as the participants were able to recognize the approximate shape and zoom quickly only on that part of the image. However, as the targets became too small to even approximately identify, virtual navigation became essential and performance (i.e. success rate) declined rapidly.

While the results show a slight advantage to the TDW for the target searches, it was not as significant as expected. This is in part because the experiment deliberately spanned a range of difficulty, and thereby is inclusive of both the SDD and TDW advantages. However, results obtained from the post experiment survey indicate that participants decidedly preferred the TDW experience over the SDD, even if their performance results did not show a marked difference. Indeed, several participants believed themselves to have




Figure 13: Ease of use of the display environments. Results were obtained from post-experiment surveys completed by all 57 participants.




Figure 14: Suitability of the display environments for target searching in ultra-high resolution images. Results were obtained from post-experiment surveys completed by all 57 participants.

performed better with the TDW when they had in fact performed better with the SDD.

The novelty of the TDW environment cannot be ignored and the fact that observing very large images in such an immersive environment will have had some emotional impact on the way participants viewed the experiment. Also, the experiment was clearly investigating the perceived value of TDWs compared to the SDD, and may have skewed participants' perception in favour of the more novel technology. A cross-section of participants primarily sourced from universities may reflect that preference for new technology. However, the primary use of the TDW is in this sector and therefore the performance of such a cohort remains relevant.

The very fact that participants preferred the TDW environment is important even when it did not correspond to increased performance. This suggests that participants might be more inclined to persist with the

TDW environment further than with the SDD, however this may be a result of the novelty factor.

No matter how much the novelty factor plays a part, we find that the experience of the participants in this study reflect the results of previous experiments that show performance improvement when virtual navigation can be avoided.

The results from this study indicate considerable opportunities for further work in general testing of TDWs and domain specific testing. This experiment used specifically constructed conditions to examine aspects of the display environments, however, these conditions aren't necessarily indicative of typical astronomical activities. Therefore, before astronomers are likely to include a TDW into their workflow, there needs to be evidence that TDW will actually improve efficiency and/or reduce errors or omissions.

Based on our understanding of how participants




Figure 15: Impact of screen elements on the task based on post-experiment survey. The bars indicate the percentage of non-astronomers (green) and astronomers (red) who provided an affirmative response to questions regarding the level of distraction caused by screen elements, and their perceived impact on the search process.

used the displays, and in response to individual comments, we identify several areas for future investigation:

- **Multiple image inspection.** Our experiments focused on one particular use case for TDWs - inspection of individual ultra-high resolution images. However, there is an alternative way to take advantage of the display pixels: instantaneous display of many individual lower-resolution images. Consider the simple case of classifying structures in images: a crucial skill in many fields of astronomy which remains difficult to fully automate. Using a TDW would allow astronomers to maintain a view of many classified structures, thereby assisting the evaluation of unclassified structures. This would provide an opportunity for refinement and reclassification, as previous decisions are still available for scrutiny.
- **Extended exposure.** The experiments in this study were designed to run within a 30-minute period to make it easier for participants to commit their time. The results from the survey indicate that most participants felt the TDW was easier to use than the SDD, as well as being more suitable to the task. This suggests that extending the exposure time in a variety of ways might yield more distinguishing results. For example, if the experiment were not time limited, how long would it take to find all targets? Alternatively, if the participants were to use the TDW each day for a period of time, for example, a week, would they improve their performance compare to a control group using SDDs?

- **Features retained over extended periods.** Following on from the previous item, it would be useful to learn whether a TDW aids in the recall of multi-scale features in ultra-high resolution images. Such an experiment would test a participant's ability to relocate features in an image that they had previously been able to find, or had been shown to them. This experiment would look at the difference between short and mid-term memory to see if the TDW exposure shows a difference compared to SDD exposure.
- **Collaborative exploration.** Our study looked at a very basic form of collaboration with two participants working together to share the task. However, there are several variations that would be worthy of further investigation. For example, rather than sharing a SDD, can the frustration caused by sharing control be alleviated by providing each participant with a SDD? It would be expected that more overlap of searched area would occur, but this might be mitigated if each participant could observe the other's display. Moreover, how exactly does communication between participants occur in this situation, either naturally or guided? Is physical proximity necessary, or is virtual proximating via teleconferencing facilities sufficient? Such a study has added relevance for the case where participants are working off physically remote but linked TDWs, as in the case of OptIPortals. Finally, increasing the number of participants beyond two, might establish a relationship between screen size and practical use with respect to the number of people observing that data.
- **Consumer 4K UHD displays.** With the recent availability of consumer-grade 4K UHD displays, it would be valuable to repeat the experiment in an environment that might represent an effective combination of the SDD and TDW. While not providing the number of pixels available on a TDW, the advantages of a SDD would be brought to bear, and might produce a cost-effective compromise. Depending on the screen size, this might also prove to be a viable collaborative environment as well as being suitable for an individual. While software like SAGE would work with a 4K display, a significant benefit of running a display from a single computer is that windowing environments can be configured easily and no inter-machine synchronisation is required. This means that all applications can be run without modification.

## 5 Conclusion

The amount of information captured by current and future astronomical instruments greatly outstrips the resolution of both current and on-the-horizon displays. TDWs provide a cost-effective method of achieving an order of magnitude increase in display resolution, thereby enhancing the presentation of astronomical data



and potentially optimising the consumption of information. However, the notion that TDWs are essential when dealing with extremely large images is not so clear.

The results from this study indicate that TDWs provide a better platform for searching for discrete targets within large images than with a SDD. It also shows that astronomers perform somewhat better than non-astronomers at extracting information from extremely large images (likely due to their more systematic approach to searching), and that the collaborative combination of two non-astronomers using a TDW rivals the experience of an individual astronomer. However, the study also indicates that there is a great variety of aptitude of participants, suggesting that TDWs might greatly enhance the performance for some individuals, while providing little help to others. It also shows that the type of content also has a significant impact on the participants ability to identify targets.

This experiment has borne out the results of earlier research highlighting the benefits of physical versus virtual navigation, and the value of TDWs for searching for very small targets. However, in astronomy, as in many other disciplines, there is a great variation in the physical scale and the “visual connectivity” of the objects to be studied. Our experiments highlight the differences between looking at words, which are processed differently by the brain, compared to identifying isolated galaxies in wide field surveys or visually continuous nebula.

A TDW provides a very impressive environment to examine images and participants enjoyed the experience, which significantly influences their perception of the suitability of the TDW environment. While such value is difficult to quantify, it suggests that the availability of a TDW can be a useful addition to the astronomer’s work flow - if only because using one is a more enjoyable task than being seated at the desktop. We are encouraged to believe that the TDW has now come-of-age for astronomy, particularly as a collaborative environment. Ultimately, the practicality of the wider up-take of TDWs for astronomy is contingent on increasing ease-of-use (e.g. through the SAGE environment), suitable interface options and simple training in the use of physical navigation.

## Acknowledgments

We thank the participants for contributing their time to this project. All experimental work was approved and conducted in accordance with the requirements of Swinburne Universitys Human Research Ethics Committee (SUHREC). We also thank Dr Simon Cropper and Dr David O’Conner (Psychology dept., University of Melbourne) for their advice on the experiment design, and Christoph Willing (University of Queensland) and Luc Renambot (University of Chicago) for their technical advice in the use of the SAGE environment. We also thank ITS Research Services at the University of Melbourne for the use of the OzIPortal for this experiment.

## References

- ACSWeb: ACS::The Advanced Camera for Surveys, <http://acs.pha.jhu.edu/> (accessed 26 September 2012)
- AdlerWeb: Shoot for the Moon – Adler Planetarium, <http://www.necdisplay.com/case-study/adler-planetarium/21> (accessed 18 July 2014)
- Andrews, C., Endert, A., & North, C. 2010, in (ACM Press), 55
- Andrews, C., Endert, A., Yost, B., & North, C. 2011, *Inf. Vis.*, 10(4), p.341–355
- AnglikWeb: Anglik, <http://www.anglik.net/english250.htm> (accessed 18 October 2013)
- Ball, N. M., & Brunner, R. J. 2010, *Int. J. Mod. Phys. D*, 19(07), p.10491106
- Ball, R., & North, C. 2005a, in (ACM Press), 1196
- Ball, R., & North, C. 2005b, in *Human-Computer Interaction - INTERACT 2005*, , ed M. F. Costabile & F. Patern (Berlin, Heidelberg: Springer Berlin Heidelberg), 350363
- Ball, R., North, C., & Bowman, D. A. 2007, in (ACM Press), 191
- Bezerianos, A., & Isenberg, P. 2012, *IEEE Trans. Vis. Comput. Graph.*, 18(12)
- Bi, X., & Balakrishnan, R. 2009, in (ACM Press), 1005
- Brescia, M., Cavuoti, S., Djorgovski, G. S., Donalek, C., Longo, G., & Paolillo, M. 2012, in *Astrostatistics and Data Mining*, , ed L. M. Sarro, L. Eyer, W. OMullane, & J. De Ridder, 31
- Calit2Web: Calit2: Calit2 and University of Melbourne Initiate Australias Ultra-Resolution Global Collaboration Laboratory, <http://www.calit2.net/newsroom/release.php?id=1219> (accessed 24 September 2012)
- CGLXWeb: CGLX Project, <http://vis.ucsd.edu/~cglx/> (accessed 26 September 2012)
- Cruz-Neira, C., Sandin, D. J., DeFanti, T. A., Kenyon, R. V., & Hart, J. C. 1992, *Commun. ACM*, 35(6), p.6472
- DarkEnergySurveyWeb: The Dark Energy Survey, <http://www.darkenergysurvey.org/> (accessed 24 September 2012)
- DeFanti, T. A. et al. 2009, *Future Gener. Comput. Syst.*, 25(2), p.114123
- DeFanti, T. A. et al. 2010, *Cent. Eur. J. Eng.*, 1(1), p.1637
- Febretti, A. et al. 2013, in *Proc. SPIE 8649*, , ed M. Dolinsky & I. E. McDowall, 86490386490312

- Fujiwara, Y., Date, S., Ichikawa, K., & Takemura, H. 2011, *J. Inf. Process. Syst.*, 7(4), p.581594
- Hassan, A., & Fluke, C. J. 2011, *Publ. Astron. Soc. Aust.*, 28(2), p.150170
- HiperwallWeb: HiperWall — hiperwall.calit2.uci.edu, <http://hiperwall.calit2.uci.edu/> (accessed 26 September 2012)
- HyperSuprimeCamWeb: Hyper Suprime-Cam, <http://www.naoj.org/Projects/HSC/index.html> (accessed 26 September 2012)
- Ivezic, Z. et al. 2008, *Serbian Astron. J.*, (176), p.113
- Keller, S. C. et al. 2007, *Publ. Astron. Soc. Aust.*, 24(1), p.1
- Koribalski, B. S. 2012, *Publ. Astron. Soc. Aust.*
- LeedsPowerwallWeb: Leeds Virtual Reality Powerwall project wins a large extension - Faculty of Engineering - University of Leeds, <http://www.engineering.leeds.ac.uk/faculty/news/2009/powerwall.shtml> (accessed 26 September 2012)
- Mohr, J. J. et al. 2012, *ArXiv Prepr. ArXiv12073189*
- Moreland, K. 2012, in *Proceedings of the IEEE Symposium on Large-Scale Data Analysis and Visualization*
- Morikawa, Y. et al. 2010, *AGU Fall Meet. Abstr.*, p.D5
- NICTWeb: Universal Communication Research Institute — Information Services Platform Laboratory — NICT-National Institute of Information and Communications Technology, <http://www.nict.go.jp/en/univ-com/isp/topics.html> (accessed 27 September 2012)
- OptiportalWeb: Optiportals, [http://optiportal.org/index.php/Main\\_Page](http://optiportal.org/index.php/Main_Page) (accessed 11 March 2014)
- OziPortalNewsWeb: Fast HD ultra-broadband link is powerful collaboration tool?: News?: The University of Melbourne, <http://archive.unimelb.edu.au/news/4894/index.html> (accessed 27 September 2012)
- Ponto, K., Doerr, K., Wypych, T., Kooker, J., & Kuester, F. 2011, *Future Gener. Comput. Syst.*, 27(6), p.649656
- SAGEWeb: SAGE: SCALABLE ADAPTIVE GRAPHICS ENVIRONMENT, <http://www.sagecommons.org/> (accessed 26 September 2012)
- Sims, M. H., Dodson, K. E., & Edwards, L. J. 2010, *LPI Contrib.*, 1538, p.5614
- Smarr, L., Brown, M., & de Laat, C. 2009, *Future Gener. Comput. Syst.*, 25(2), p.109113
- Tada, T., Shimojo, S., Ichikawa, K., Abe, H., & others 2011, in *Granular Computing (GrC)*, 2011 IEEE International Conference on, 613618
- Taesombut, N. et al. 2006, *Future Gener. Comput. Syst.*, 22(8), p.955963
- Tyson, J. A. 2002, in *Proceedings of the SPIE*, , ed J. A. Tyson & S. Wolff, 1020
- Way, M. J., Scargle, J. D., Ali, K. M., & Srivastava, A. N. 2012, *Advances in Machine Learning and Data Mining for Astronomy* (CRC Press)
- Yamaoka, S., Manovich, L., Douglass, J., & Kuester, F. 2011, *Computer*, 44(12), p.3948
- Yost, B., Haciahetoglu, Y., & North, C. 2007, *ACM CHI*, p.101111

## A The OzIPortal Project

We now look at the OptIPortal project and a specific TDW, the OzIPortal, in more detail in order to understand some of the reasons why these devices have not already become standard elements of the research workflow across diverse scientific disciplines.

### A.1 The OptIPortal Project

The OptIPortal project grew out of the Optiputer project, a US government funded project to connect high performance computing facilities together via optical networks, called Lambdas (DeFanti et al. 2009; Taesombut et al. 2006). With such powerful data processing and transfer resources in the background, an enhanced visualization capability was required. This was initially achieved using several projectors, with edge-blending techniques to compensate for luminosity fall-off between adjacent projections. The OptIPortal project took this methodology to the next level, adopting high-resolution off-the-shelf displays to create tiled surfaces. Software was developed to make the management of the TDW relatively transparent so that the users could focus on the research content (DeFanti et al. 2009).

Collaborative environments such as the OptIPortal network, were designed to allow distributed research teams to work together simultaneously on the resulting imagery and analysis. This would provide greater opportunity for collaborative research, leading to greater understanding of the data (Smarr et al. 2009; Sims et al. 2010; Yamaoka et al. 2011).

While the project promised a simple, powerful and interconnected system, there were many problems with the early incarnations of OptIPortals, primarily due to immaturity of the associated software. To tackle this, workarounds are commonly employed, such as manual data processing steps. Whilst giving the appearance of success for visualization, this led to some misunderstandings as to what an OptIPortal actually is and its overall utility as part of research workflows.

## A.2 The OzIPortal Experience

In 2008, The University of Melbourne launched the OzIPortal, a 98 megapixel TDW, with considerable fanfare. It was lauded as an amazing research tool: "In an Australian first, this next-generation platform set to revolutionize the way Australia interacts with the rest of the world allows real-time, interactive collaboration across the globe, combining high-definition video and audio with the sharing of ultra-resolution visualizations from a broad range of disciplines." (Calit2Web 2012).

Despite the high level of interest generated by early demonstrations, the OzIPortal failed to attract a significant commitment from the research community.

The OzIPortal was configured using 24 x 2560x1600 LCD displays, initially in an 8x3 arrangement and later in 6x4. 12 slave nodes with dual-head graphics cards were used to drive two monitors each. Reducing the number of slave nodes improved the operation of the TDW without reducing the performance. Data was exported from the head node via NFS to each of the slave nodes. An additional machine was required to provide a real-time video stream that would allow any video signal captured via HDMI to be presented on the TDW. In this way, the OzIPortal was able to include low latency, high-definition video conferencing content alongside other stored content.

One of the initial drivers for the OzIPortal was to establish a dedicated gigabit link from The University of Melbourne through to CALIT2 in the United States (OziPortalNewsWeb 2012). This was successfully implemented on a layer 2 network via AARNet. The purpose of the network was to demonstrate the rapid transfer of massive datasets, to show how distant collaborators could work on the same datasets at the same time. In practice however, it was necessary to distribute the data in advance, as some of the data was too large to deliver in a timely manner, even over the dedicated link. Instead, the link was used primarily to stream uncompressed video directly to the TDW, instead of via a video conference codec. Finally, it took several technicians to satisfactorily operate the TDW, and a great deal of testing beforehand was required to minimize disruptions during events. Day-to-day operations could not be sustained with such human resource demands and as such provided a less than satisfactory experience for users.

Ultimately, what was required for the successful deployment of a TDW in the research workflow was a more stable system, that was easier to use, did not require high-level of support, automated preparation of content, and the ability to run applications specific to individual scientific disciplines. Our recent experiences with the OzIPortal, particularly through the use of the SAGE environment, is a positive step towards these outcomes.